

# Averaged Hausdorff Approximations of Pareto Fronts based on Multiobjective Estimation of Distribution Algorithms

2015

Luis Martí

Departamento de Engenharia Elétrica  
PUC Rio de Janeiro, Brazil  
lmarti@ele.puc-rio.br

Christian Grimme

Information Systems and Statistics  
University of Münster, Germany  
christian.grimme@wi.uni-muenster.de

Pascal Kerschke

Information Systems and Statistics  
University of Münster, Germany  
kerschke@uni-muenster.de

Heike Trautmann

Information Systems and Statistics  
University of Münster, Germany  
trautmann@wi.uni-muenster.de

Günter Rudolph

Department of Computer Science  
TU Dortmund University, Germany  
guenter.rudolph@tu-dortmund.de

## Abstract

In the a posteriori approach of multiobjective optimization the Pareto front is approximated by a finite set of solutions in the objective space. The quality of the approximation can be measured by different indicators that take into account the approximation's closeness to the Pareto front and its distribution along the Pareto front. In particular, the averaged Hausdorff indicator prefers an almost uniform distribution. An observed drawback of multiobjective estimation of distribution algorithms (MEDAs) is that - as common for randomized metaheuristics - the final population usually is not uniformly distributed along the Pareto front. Therefore, we propose a postprocessing strategy which consists of applying the averaged Hausdorff indicator to the complete archive of generated solutions after optimization in order to select a uniformly distributed subset of nondominated solutions from the archive. In this paper, we put forward a strategy for extracting the above described subset. The effectiveness of the proposal is contrasted in a series of experiments that involve different MEDAs and filtering techniques.

# 1 Introduction

Many real-world optimization problems involve more than one goal to be optimized and are known as multiobjective optimization problems (MOPs), i.e. a set of objective functions  $f_1(\mathbf{x}), \dots, f_M(\mathbf{x})$  should be jointly optimized; formally,

$$\min \mathbf{F}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_M(\mathbf{x})); \mathbf{x} \in \mathcal{S}; \quad (1)$$

where  $\mathcal{S} \subseteq \mathbb{R}^n$  is known as the feasible set and could be expressed as a set of restrictions over the decision or search space  $\mathbb{R}^n$ . The image set  $\mathcal{O} \subseteq \mathbb{R}^M$  of  $\mathcal{S}$  produced by the vector-valued function  $\mathbf{F}(\cdot)$  is called feasible objective set or criterion set. The solution to this type of problem is a set of trade-off points. The optimality of a solution can be expressed in terms of the Pareto dominance relation.

**Definition 1 (Pareto dominance)** *In the optimization problem (1) and having  $\mathbf{x}, \mathbf{y} \in \mathcal{S}$ ,  $\mathbf{x}$  is said to dominate  $\mathbf{y}$  (expressed as  $\mathbf{x} \prec \mathbf{y}$ ) if  $\forall j = 1, \dots, M: f_j(\mathbf{x}) \leq f_j(\mathbf{y})$  and  $\exists i \in \{1, \dots, M\}: f_i(\mathbf{x}) < f_i(\mathbf{y})$ . The non-dominated subset  $\mathcal{A}^*$  of set  $\mathcal{A} \subseteq \mathcal{S}$  is defined as*

$$\mathcal{A}^* = \{\mathbf{x} \in \mathcal{A} \mid \nexists \mathbf{x}' \in \mathcal{A} : \mathbf{x}' \prec \mathbf{x}\}.$$

The solution of (1) is  $\mathcal{S}^*$ , the non-dominated subset of  $\mathcal{S}$ .  $\mathcal{S}^*$  is known as the *efficient set* or *Pareto-optimal set* [1]. Its image in objective space is known as the *Pareto-optimal front*,  $\mathcal{O}^*$ . As finding the explicit formulation of  $\mathcal{S}^*$  is often impossible, generally, an algorithm solving (1) yields a discrete non-dominated set,  $\mathcal{P}^*$ , that approximates  $\mathcal{S}^*$ . The image of  $\mathcal{P}^*$  in objective set,  $\mathcal{PF}^*$ , is known as the *non-dominated front*.

A broad range of heuristic and metaheuristic approaches has been used to address MOPs [1]. Of these, evolutionary multiobjective optimization algorithms (EMOAs) [2] have been found to be a competent approach in a wide variety of application domains. Alternatively, multiobjective estimation of distribution algorithms (MEDAs) were introduced which aim at learning the problem structure and characteristics along the run. In this paper, we will experimentally investigate how representative MEDAs perform compared to classical EMOAs and how especially MEDAs can be improved after the run by means of a postprocessing approach in order to get more equally spaced solutions on the final Pareto front approximation.

The crucial task is how to measure the performance in the multiobjective setting, i.e. how to assess the relation of  $\mathcal{PF}^*$  to  $\mathcal{O}^*$ . Several performance indicators have been proposed including the hypervolume indicator or the R indicators, see [3, 4] for an overview. Each indicator concentrates on special desired characteristics of the front approximation while one frequently discussed aim is that elements of  $\mathcal{PF}^*$  should be evenly spread along the true Pareto front in order to present an unbiased solution set to the decision maker (e.g. [5]). On the application side, specifically optimization systems or online control could profit from such approximations such that changes from one solution to the other are quite moderate and of equal extent. In [6] an approach for evolving Pareto front approximations with uniform gap is presented. However, most indicators do not address this issue.

The  $\Delta_p$  indicator introduced in [7] specifically favors approximations with the desired characteristic based on averaged Hausdorff distances between  $\mathcal{PF}^*$  and  $\mathcal{O}^*$ . Several algorithmic approaches have been introduced so far (e.g. [8, 9, 10]) to generate quite equally spaced Pareto front approximations but all of them are focusing on integrating the  $\Delta_p$  based subset selection into the EMOA. In this paper, we present a strategy to be applied posterior to an approximation run to extract a most uniformly distributed subset from the archive of all nondominated solutions generated within the course of the used EMOA.

The remainder of this paper is organized as follows: In Section 2 the methodological background is provided regarding multiobjective estimation of distribution algorithms and the postprocessing strategy. Experimental results are presented and discussed in Section 3. Conclusions are drawn in Section 4 supplemented by an outlook on further research directions.

## 2 Methodology

### 2.1 Multiobjective Estimation of Distribution Algorithms

The inclusion of learning as part of the search process can improve the performance of evolutionary multiobjective optimization algorithms [11]. It can be argued that learning would allow to grasp the characteristics of the problem being solved and, hence, explore the search space in a more efficient manner. There are some approaches that perform this task by providing hybrid evolutionary/machine learning methods, like, for example, the learnable evolution model (LEM) [12]. Other approaches attempt to infer the fitness landscape of the problem in order to find promising search directions [13, 14, 15]. Another alternative for carrying out this task is to resort to what has been denominated as estimation of distribution algorithms (EDAs) [16]. This is because of EDAs capacity of learning the problem structure. EDAs replace the application of evolutionary operators in the offspring generation process with the creation of a statistical model of the fittest elements of the population in a process known as model-building. This model is then sampled to produce new elements. Nevertheless, the so called multiobjective EDAs (MEDAs) [17] have not lived up to their a priori expectations. This can be attributed to the fact that most MEDAs have limited themselves to transforming single-objective EDAs into a multiobjective formulation by including an existing multiobjective fitness assignment function.

Bosman and Thierens [18] proposed the successful multiobjective mixture-based iterated density estimation algorithms (MIDEAs). They also proposed a novel Pareto-based and diversity-preserving fitness assignment function. MIDEA considered several types of probabilistic models for both discrete and continuous problems. A mixture of univariate distributions and a mixture of tree distributions were used for discrete variables. A mixture of univariate Gaussian models and a mixture of multivariate Gaussian factorizations were applied for continuous variables. Different variants of MIDEAs implement this in a particular form. For example, naive MIDEA the naive Bayes algorithm.

It can also be argued that the covariance matrix adaptation evolution strategies (CMA-ES) [19] are also EDAs, as they construct an abstract model of the population and then sample it to produce new individuals. CMA-ES provides a method for updating the covariance matrix of the multivariate normal mutation distribution used in an evolution strategy [20]. The covariance matrix describes the pairwise dependencies between the variables in the distribution. Adaptation of the covariance matrix is equivalent to learning a second-order model of the underlying objective function. The multi-objective CMA-ES (MO-CMA-ES) [21] is the extrapolation to the multi-objective domain.

It has been pointed out that the current model-building algorithms of MEDAs have a set of drawbacks that would prevent those algorithms from yielding substantially better results. In particular, the tendency of MEDAs loosing population diversity has repeatedly been reported [22, 23]. This situation, although already described in single-objective EDAs [24], is particularly dramatic in the multiobjective case, as diversity and homogeneity are among the desired features of the final non-dominated set. An analysis of the results yielded by current MEDAs and their scalability with regard to the number of objectives leads to hypotheses regarding the issues that might be hampering the obtention of substantially better results with regard to other evolutionary approaches [25].

The incorrect handling of data outliers is a paradigmatic example of insufficient comprehension of the nature of the model building problem. In ‘standard’ machine learning practice, outliers are treated as noisy, inconsistent or irrelevant data. Therefore, this type of data is expected to have little influence on the model or just to be disregarded. However, that behavior is not adequate for model building. In this case, it is known beforehand that all elements in the data set should be taken into account as they represent newly discovered or candidate regions of the search space and therefore must be explored. Therefore, these instances should be at least equally represented by the model and perhaps even reinforced. This situation is further aggravated by the mating selection scheme employed in most MEDAs. The continuous selection of the best part of the

population might lead to a premature homogenization of the population and, therefore, to the stagnation of the search process. In the second case, the loss of diversity can be traced back to the above-described outliers issue of model-building algorithms and also to the incorrect estimation or sampling of the model. This fact leads us back to the statement referring that model building has not been correctly acknowledged as a different problem with particular requirements.

There have been some works that propose to improve this situation by introducing model building algorithms better suited for the task. That is the case of the MIDEA algorithm family, with the introduction of the adaptive variance scaling (AVS) and the standard deviation ratio (SDR) [26]. The AVS and SDR combination helps fight the early reduction of the mixture densities variances and therefore the premature convergence and diversity loss. Another important milestone has been the introduction of the anticipated mean shift (AMS) that takes into account the previous values of the means of the distribution to “push” solutions towards the Pareto-optimal front. AMS has been conjointly used with AVS in the multiobjective adapted maximum-likelihood Gaussian mixture model (MAMaLGaM-X) [27]. Similarly, the multiobjective optimization neural EDA (MONEDA) [28] embeds a custom-made model building algorithm [29] that is able to maintain diversity by correctly handling the outlier elements. This approach has been improved by the introduction of the match-based learning paradigm of adaptive resonance theory (ART) [30] leading to the multiobjective ART EDA (MARTEDA) [31].

Other approaches propose modifications to currently existing methods. For example, in [23] a method that avoids too sudden drops (to zero) of the variances of a multivariate Gaussian model is proposed. Similarly, in [32] a permutation sampling is introduced that eliminates the sampling errors of UMDA [33]. Other approaches [34, 35] have tried strategies to re-inject “fresh” individuals that are kept on a evolutionary algorithm that is run in parallel. In spite of these efforts, the issue of population diversity and lack of homogeneity is still an open topic that must be addressed.

## 2.2 Postprocessing of MEDA results

Suppose we are particularly interested in a finite size approximation of  $\mathcal{O}^* = \mathbf{F}(\mathcal{S}^*)$  that exhibits a sufficiently good spread and a small distance to  $\mathcal{O}^*$ . Previous work in the context of EMOA (see e.g. [9]) has shown that the *averaged Hausdorff distance* [7] can be used to achieve this goal.

**Definition 2** Let  $A, B \subset \mathbb{R}^M$  be non-empty finite sets. The value

$$\Delta_p(A, B) = \max(\text{GD}_p(A, B), \text{IGD}_p(A, B)) \text{ with}$$

$$\text{GD}_p(A, B) = \left( \frac{1}{|A|} \sum_{a \in A} d(a, B)^p \right)^{1/p} \text{ and}$$

$$\text{IGD}_p(A, B) = \left( \frac{1}{|B|} \sum_{b \in B} d(b, A)^p \right)^{1/p}$$

for  $p > 0$  is termed the averaged Hausdorff distance between sets  $A$  and  $B$  where  $d(u, A) := \inf\{\|u - v\| : v \in A\}$  for  $u, v \in \mathbb{R}^M$  and some vector norm  $\|\cdot\|$ .

Suppose that some MEDA has generated an approximation of the Pareto front for some MOP. Typically, this approximation does not yield a finite point set in objective space that is evenly distributed. Therefore, we propose the following *postprocessing* approach:

1. Run your favorite MEDA that is equipped with a tiny add-on: store a copy of each offspring that is generated and evaluated in a file.
2. After termination of your favorite MEDA:
  - construct an evenly spaced reference front from a given approximation of the Pareto front (e.g. from last population of favorite MEDA); then feed each stored offspring into the  $\Delta_p$  archive updater (see alg. 1) sequentially; the final content of the archive  $A$  is the desired approximation.

For bi-objective problems the reference front  $R$  can be constructed as follows: calculate a linear interpolation from a given approximation of the Pareto front. Since the length  $L$  of the resulting polygonal line can be derived easily, division by the number of desired archive points  $m$  yields the size of the equal spacing  $\delta = L/m$  which is used to place  $m$  points along the polygonal line with distance  $\delta$ , where the extreme points of the discretization are moved half the length inwards the polygonal line. After the reference front  $R$  has been constructed it is used in the  $\Delta_p$  archive updater to decide which point should be added to or deleted from the archive.

An update operation can be realized as sketched in alg.1. This naive approach takes  $\Theta(|A| \cdot (|A| \cdot |R| \cdot M))$  time units, but there exists a quick update version [9] that only needs  $\Theta(|A| \cdot |R| \cdot M)$  time units.

---

**Algorithm 1** Naive  $\Delta_1$ -update [9]

---

**Require:** archive set  $A$ , reference set  $R$ , new element  $\mathbf{x}$

```

1:  $A = \text{ND}_f(A \cup \{\mathbf{x}\}, \preceq)$ 
2: if  $|A| > N_R := |R|$  then
3:   for all  $a \in A$  do
4:      $h(a) = \Delta_1(A \setminus \{a\}, R)$ 
5:   end for
6:    $A^* = \{a^* \in A : a^* = \text{argmin}\{h(a) : a \in A\}\}$ 
7:   if  $|A^*| > 1$  then
8:      $a^* = \text{argmin}\{\text{GD}_p(A \setminus \{a\}, R) : a \in A^*\}$ 
9:   end if
10:   $A = A \setminus \{a^*\}$ 
11: end if
```

---

The most obvious order of feeding the stored pairs  $(\mathbf{x}, \mathbf{F}(\mathbf{x}))$  into the archive updater is the order of their generation. We call this the ‘forward update.’ In this manner, many individuals will pass the initial dominance check, so that subsequent  $\Delta_p$  calculations are necessary. Some time saving may be achieved by feeding the stored pairs into the archive update in inverted order. We call this the ‘backward update.’ Since points that have been generated in later iterations of the MEDA are more likely to dominate previous points, most points from the rear of the inverted sequence will probably not pass the initial dominance check, so that subsequent  $\Delta_p$  calculations can be avoided. Since the order of the points presented to the archive clearly affects the final outcome of the archive, we shall compare both approaches experimentally in the next section.

## 3 Experiments

### 3.1 Experimental Setup

In order to experimentally evaluate the  $\Delta_p$  archive approximation for multiple multiobjective EDAs, we implemented two offline strategies in the JAVA<sup>TM</sup> jMetal framework [36]: a forward  $\Delta_p$ -approach ( $+fDP$ ), starting with the initial population and iteratively updating with previously traced generations, whereas the backward-strategy ( $+bDP$ ) starts with the final population. Within the  $\Delta_p$ -update mechanisms, a linear interpolation and a PSA-based strategy, both using the aggregated solutions of the whole run, were used for deriving a set of evenly spaced reference points. Based on that set, a solution was eventually added to the archive.

We then used well known MOPs for benchmarking [8]: an instance of the sphere problem with  $X \subseteq \mathbb{R}^2$  and a connected convex Pareto front, the DENT problem with  $X \subseteq \mathbb{R}^2$  and a connected convex-concave-convex front, ZDT3 with  $X \subseteq \mathbb{R}^{20}$  and a disconnected convex/concave front, as well as WFG1 with  $X \subseteq \mathbb{R}^{2+4}$ , also with a convex-concave front.

Then, reference fronts covered by 1,000 uniformly spaced points were created in order to compare our  $\Delta_p$  approximation quality for all algorithms. With the exception of WFG1, we were able to

<b>MARTEDA</b>	
F2 vigilance threshold ( $\rho$ )	0.05
Initial standard deviations ( $\varphi$ )	<b>0.01</b>
Selection percentile ( $\alpha$ )	0.3
$\hat{P}_t$ to $N^*$ ratio ( $\gamma$ )	0.5
Substitution percentile ( $\omega$ )	0.25
<b>MONEDA</b>	
Number of initial GNG nodes ( $N_0$ )	2
Maximum edge age ( $\nu_{\max}$ )	40
Best node learning rate ( $\epsilon_b$ )	0.1
Neighboring nodes learning rate ( $\epsilon_v$ )	0.05
Insertion error decrement rate ( $\delta_I$ )	0.1
General error decrement rate ( $\delta_G$ )	0.1
Accumulated error threshold ( $\rho$ )	0.2
Selection percentile ( $\alpha$ )	0.3
$\hat{P}_t$ to $N_{\max}$ ratio ( $\gamma$ )	0.5
Substitution percentile ( $\omega$ )	0.25
<b>Naive MIDEA</b>	
Selection percentile ( $\tau$ )	0.3
Diversity percentile ( $\delta$ )	15
Number of parents of a variable ( $\kappa$ )	2
Maximum number of clusters	$\lceil 0.5 \lfloor \tau n_{\text{pop}} \rfloor \rceil$
Threshold for the leader algorithm	0.1
<b>MO-CMA-ES</b>	
Initial step size ( $\sigma$ )	1
Damping for step-size ( $d$ )	$1 + n_{\text{pop}}/2$
Target success rate	0.181
Cumulation time horizon	$2/(n_{\text{pop}} + 2)$
Covariance learning rate	$2/(n_{\text{pop}}^2 + 6)$
Threshold success rate	0.44

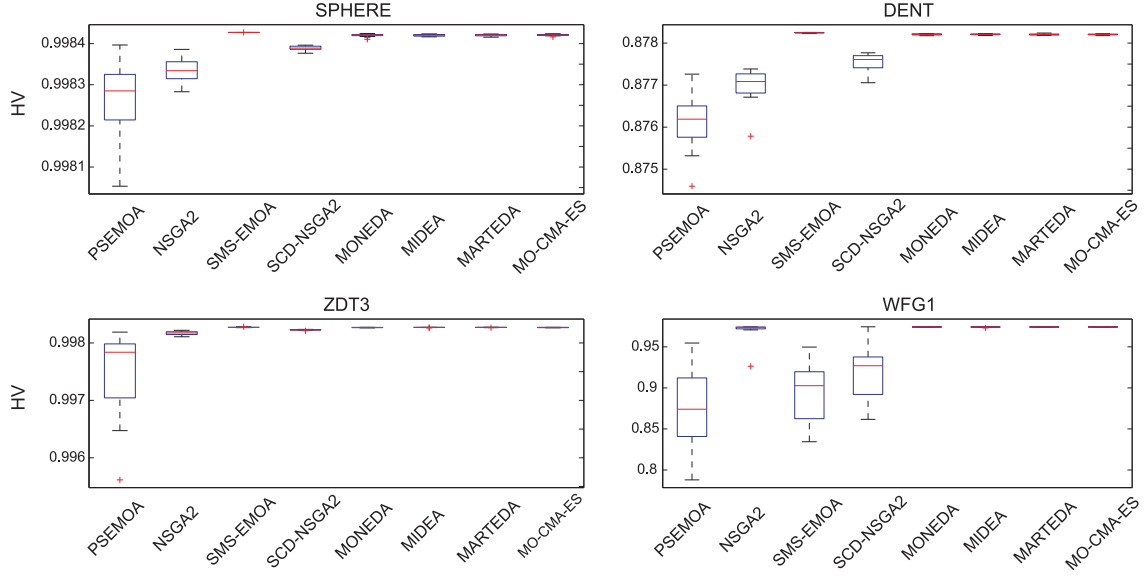
**Table 1:** Parameters of the algorithms used in the experiments following the notation of the corresponding author(s).

find a parametric form for exactly calculating the optimal fronts' length  $L$  (by rectification), for all benchmark problems. Given that form 1,000 points were placed uniformly on the PF using distance  $\Delta L = L/1,000$  between points and starting with the extreme points moved inwards by  $\Delta L/2$ . In case of WFG1 [37], the neighborhood of the known Pareto set was explored using a very fine grid (per dimension of the decision space), in order to find a PF. Then 1,000 well-distributed solutions among the non-dominated solutions were kept. As the number of reference solutions is by far larger than the approximated set's size, a not perfectly uniform distribution of the Pareto-optimal solutions is acceptable.

Our evaluation is based on four state-of-the-art general purpose EMOAs and four MEDAs. From the former group, we employed the NSGA2 [5] and SMS-EMOA [38] with standard parameter settings – SBX crossover with  $p_x = 0.9$  and polynomial mutation with  $p_{\text{mut}} = 1/n$  – as well as PSEMOA as special purpose EMOA with inherent internal clustering for generating evenly distributed solutions. For the same reason we also used a modified version of NSGA2 [39] – the *sequential crowding distance NSGA2* (SCD-NSGA2) – which successively alternates between removing the individuals with the smallest crowding distance and recomputing the crowding distance values, until only  $\mu$  solutions are left.

Considering the MEDAs we deal with two well-known approaches: the naive MIDEA [18] and the multiobjective CMA-ES (MO-CMA-ES) [21]. We also include MONEDA [28] and MARTEDA [31] as they are supposed to have a better handling of diversity. The parameters of the MEDAs are summarized in Table 1.

During our experiments, all test problems were optimized 20 times by all algorithms, given a budget of 50,000 function evaluations each and population sizes  $\mu \in \{10, 20, 100\}$ . The data, which were generated from the optimizers' runs, have been submitted to the offline  $\Delta_p$  archivers. Both update strategies were parameterized with  $p \in \{1, 2\}$ .



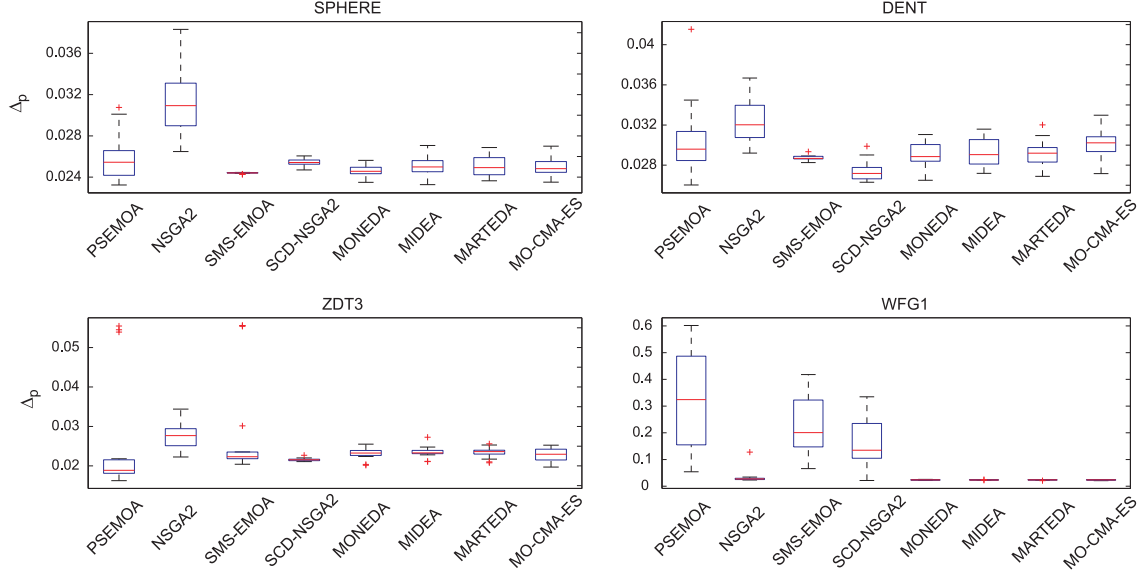
**Figure 1:** Comparison of MEDAs with EMOAs regarding Hypervolume.

### 3.2 Results

At first, overall algorithm behavior is investigated by comparing the MEDAs with the classical EMOA approaches in terms of the Dominated Hypervolume indicator (HV, [4]) as well as  $\Delta_p$  (Figures 1 and 2) based on the generated reference fronts as described in Section 3.1. It becomes obvious that the MEDAs lead to very stable results over the repeated runs which is reflected by both indicators while no MEDA is superior to the others. Moreover, for all test problems the respective performance is comparable to the best performing algorithm out of the classical EMOA set. While over all test problems but WFG1 PSEMOA and SCD-NSGA2 can be considered as best regarding  $\Delta_p$ , the SMS-EMOA unsurprisingly outperforms the other classical EMOA in terms of HV as it internally optimizes the HV indicator. The test problem WFG1 results in different algorithm rankings within the classical EMOA, i.e. NSGA2 clearly is best regarding both indicators. The stability of MEDA results is very noticeable here as the variability of PSEMOA, SMS-EMOA as well as SCD-NSGA2 is large. Concluding, the MEDAs are at the least competitive with the classical EMOAs on the considered test problems.

In Figures 3 - 5 results of the postprocessing strategies applied to the MEDA variants are visualized in terms of  $\Delta_p$  regarding the generated reference fronts for population sizes  $\mu \in \{10, 20, 100\}$ . In general, experimental results become more distinguishable with increasing population size but the basic findings are reflected for the smallest population size as well. However, varying  $p$  does not have a noticeable effect in this setting, thus the analysis concentrates on  $p = 1$  for illustration purposes.

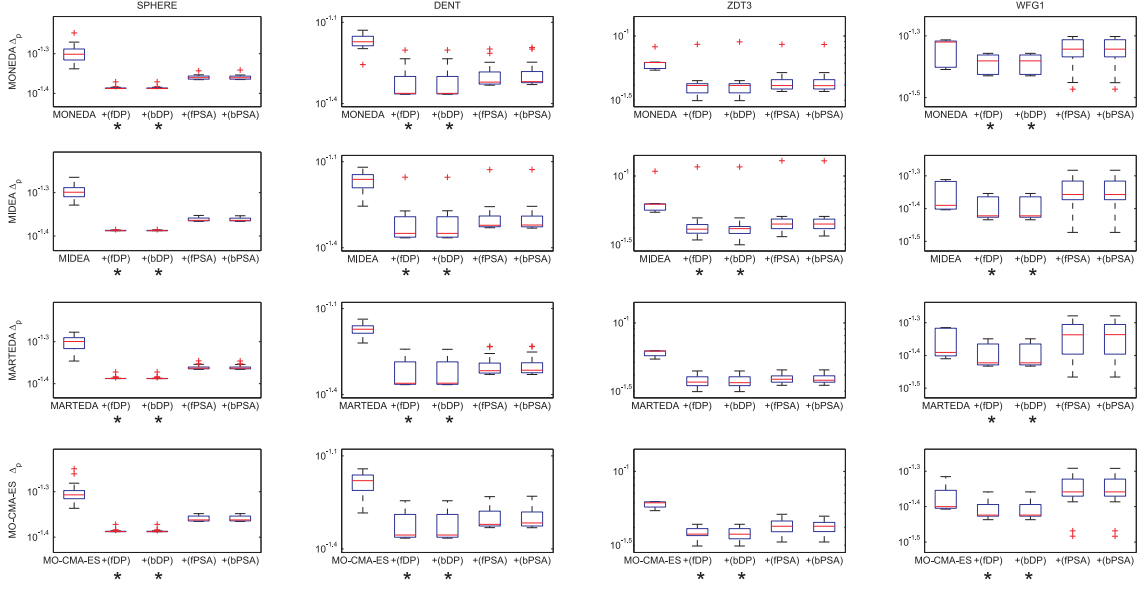
Clearly, the postprocessing improves the final approximation quality of the MEDA's run as the primary algorithm in terms of  $\Delta_p$  for all population sizes. As the Pareto front of ZDT3 is disconnected, the performance of the considered postprocessing differs from the remaining test problems. Here, the two PSA-based strategies ( $+(fPSA)$  and  $+(bPSA)$ ) outperform the archive strategy using linear interpolations of the set of all nondominated points generated in the course of the algorithm run ( $+(fDP)$  and  $+(bDP)$ ). This is intuitive as the latter approach ignores the discontinuity in the interpolation and tends to place reference points also in unattainable regions [40]. Thus, this effect increases with the population size as more points have to be placed on the linear interpolation of the front. For  $\mu = 10$  the strategies  $+(fDP)$  and  $+(bDP)$  are comparable or even slightly better than the PSA-based ones. For all other test problems  $+(fDP)$  and  $+(bDP)$  generate the largest improvement in terms of  $\Delta_p$ . Differences are statistically significant in most settings reflected by a Wilcoxon rank sum test ( $p \leq 0.05$ ) and annotated with a  $\star$  in Figures 3 - 5.



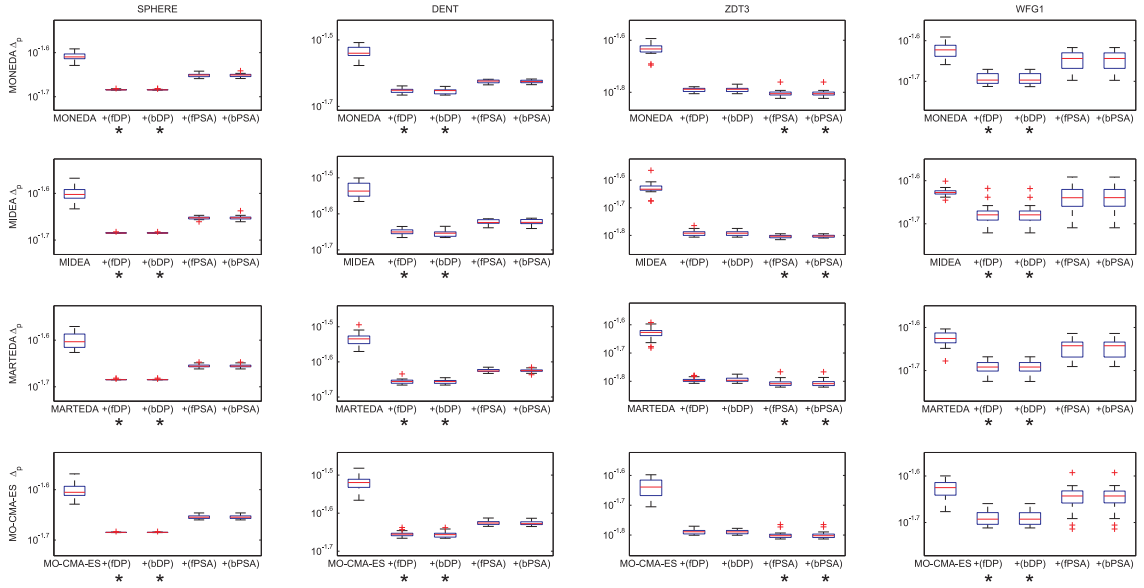
**Figure 2:** Comparison of MEDAs with EMOAs regarding  $\Delta_p$ .

Interestingly, for  $\mu = 100$  and WFG1 differences between the primary MEDA and all postprocessing variants are very small and not statistically significant. This is due to the special challenges the test problem poses onto the algorithms. Thorough analysis of the final approximation sets of the MEDAs revealed that those runs do not allow for substantial improvements of the final solution sets as the algorithm performance is not good enough and the archive of solutions does not provide the means for selecting a better subset regarding  $\Delta_p$ . However, substantially increasing the budget of the MEDAs will lead to results showing the same tendencies as described above for smaller population sizes. As expected, the direction of the archive update strategy has no influence on the final approximation quality. However, our experiments revealed that the backward strategy which starts from the final population of the primary MEDA requires much less updating iterations until the selected subset becomes stable than the forward strategy. Therefore, we recommend to favor this approach over the forward update in general such that the postprocessing will be much more computationally efficient.

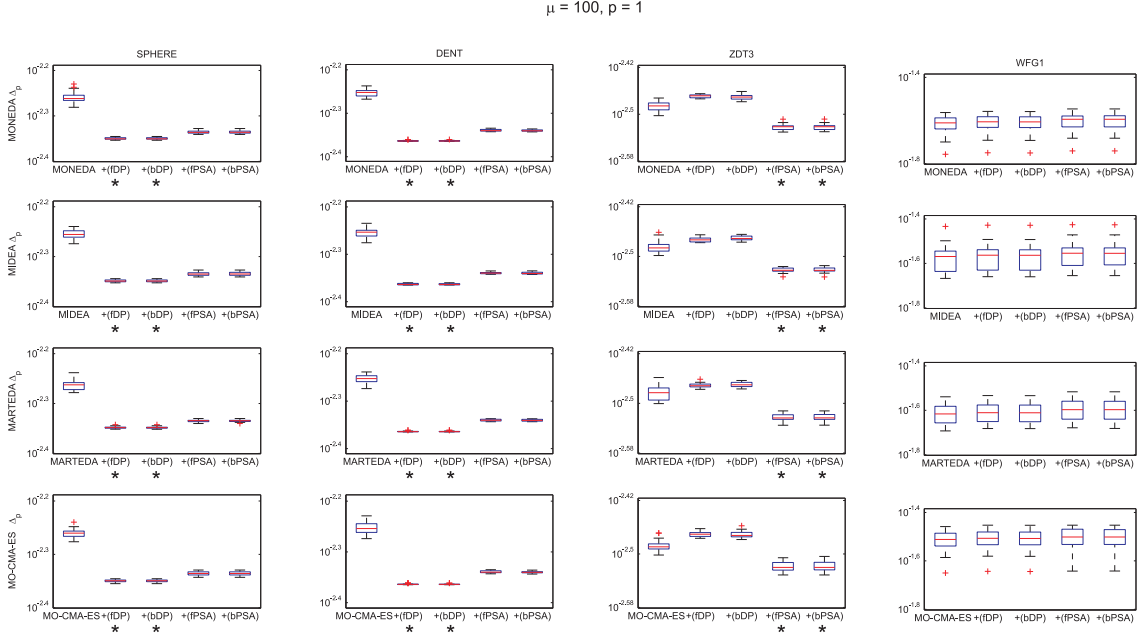


$\mu = 10, p = 1$ 

**Figure 3:**  $\Delta_p$ - and PSA-postprocessing applied to MEDA results for  $\mu = 10$  and  $p = 1$ .

 $\mu = 20, p = 1$ 

**Figure 4:**  $\Delta_p$ - and PSA-postprocessing applied to MEDA results for  $\mu = 20$  and  $p = 1$ .



**Figure 5:**  $\Delta_p$ - and PSA-postprocessing applied to MEDA results for  $\mu = 100$  and  $p = 1$ .

## 4 Conclusions and Outlook

Within this work, we first compared the quality of the solutions of four evolutionary multiobjective optimization algorithms (NSGA2, PSEMOA, SCD-NSGA2, SMS-EMOA) and four multiobjective estimation of distribution algorithms (MARTEDA, MO-CMA-ES, MONEDA, Naive MIDEA) on a set of well-known multiobjective test problems. Based on two considered indicators – HV and  $\Delta_p$  – it could be shown that the performance of the MEDAs is at least competitive to classical EMOAs. Also, the MEDAs itself performed quite similar to each other.

Due to those competitive and thus promising results, we analyzed whether the performance of the latter ones regarding  $\Delta_p$  could even be improved by a subsequent application of either forward or backward  $\Delta_p$ -update strategies which will lead to a more uniformly distributed solutions along the approximated Pareto front. Experiments confirmed this hypothesis and apparently the influence of using either one of those update directions is negligible. On the other hand, the  $\Delta_p$ -based techniques outperform the PSA-based ones besides for problems with disconnected Pareto fronts.

In future work, one might want to improve the MEDAs performance either based on the PSA- or  $\Delta_p$ -update strategy by integrating that approach into the algorithm’s mechanics. This hopefully leads to algorithms, which generate comparable or even better solutions within a single run and therefore without employing any additional postprocessing.

## References

- [1] Jürgen Branke, Kaisa Miettinen, Kalyanmoy Deb, and Roman Słowiński, editors. *Multiobjective Optimization: Interactive and Evolutionary Approaches*, volume 5252 of *Lecture Notes in Computer Science*. Springer Verlag, Berlin/Heidelberg, 2008.
- [2] Carlos A. Coello Coello, Gary B. Lamont, and David A. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Genetic and Evolutionary Computation. Springer, New York, second edition, 2007.

- [3] Eckart Zitzler, Marco Laumanns, Lothar Thiele, Carlos M. Fonseca, and Viviane Grunert da Fonseca. Why quality assessment of multiobjective optimizers is difficult. In W. B. Langdon, E. Cantú-Paz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M.A. Potter, A.C. Schultz, J.F. Miller, E. Burke, and N. Jonoska, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2002)*, pages 666–673, San Francisco, California, July 2002. Morgan Kaufmann Publishers.
- [4] Eckart Zitzler, Lothar Thiele, Marco Laumanns, Carlos M. Fonseca, and Viviane Grunert da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003.
- [5] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, April 2002.
- [6] Michael Emmerich, André Deutz, Johannes Krusselbrink, and PradyumnKumar Shukla. Cone-based hypervolume indicators: Construction, properties, and efficient computation. In RobinC. Purshouse, PeterJ. Fleming, CarlosM. Fonseca, Salvatore Greco, and Jane Shaw, editors, *Evolutionary Multi-Criterion Optimization*, volume 7811 of *Lecture Notes in Computer Science*, pages 111–127. Springer Berlin Heidelberg, 2013.
- [7] O. Schütze, X. Esquivel, A. Lara, and C. A. Coello Coello. Using the averaged Hausdorff distance as a performance measure in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 16(4):504–522, 2012.
- [8] K. Gerstl, G. Rudolph, O. Schütze, and H. Trautmann. Finding evenly spaced fronts for multiobjective control via averaging Hausdorff-measure. In *Proceedings of 8th International Conference on Electrical Engineering, Computing Science and Automatic Control (CCE)*, pages 1–6. IEEE Press, 2011.
- [9] G. Rudolph, H. Trautmann, S. Sengupta, and O. Schütze. Evenly spaced Pareto front approximations for tricriteria problems based on triangulation. In *Proceedings of 7th International Conference on Evolutionary Multi-Criterion Optimization (EMO 2013)*, pages 443–458, Berlin Heidelberg, 2013. Springer, LNCS 7811.
- [10] C. Dominguez-Medina, G. Rudolph, O. Schütze, and H. Trautmann. Evenly spaced Pareto fronts of quad-objective problems using PSA partitioning technique. In *Proceedings of 2013 IEEE Congress on Evolutionary Computation (CEC 2013)*, pages 3190–3197, Piscataway (NJ), 2013. IEEE Press.
- [11] David Wolfe Corne. Single objective = past, multiobjective = present, ??? = future. In Zbigniew Michalewicz, editor, *2008 IEEE Conference on Evolutionary Computation (CEC)*, Piscataway, New Jersey, 2008. IEEE Press.
- [12] R. S. Michalski. Learnable evolution model: Evolutionary processes guided by machine learning. *Machine Learning*, 38:9–40, 2000.
- [13] J. Deon Garrett. *Multiobjective Fitness Landscape Analysis and the Design of Effective Memetic Algorithms*. PhD thesis, Memphis State University, Memphis, TN, USA, 2008. AAI3310115.
- [14] Sébastien Verel, Laetitia Jourdan, Clarisse Dhaenens, and Arnaud Liefooghe. Set-based multi-objective fitness landscapes: definition, properties. In *4th Workshop on Theory of Randomized Search Heuristics*, Paris, France, March 2010.
- [15] Jérémie Humeau, Arnaud Liefooghe, El-Ghazali Talbi, and Sébastien Verel. Paradiseo-MO: From fitness landscape analysis to efficient local search algorithms. *Journal of Heuristics*, 19(6):881–915, 2013.

- [16] J. A. Lozano, P. Larrañaga, I. Inza, and E. Bengoetxea, editors. *Towards a New Evolutionary Computation: Advances on Estimation of Distribution Algorithms*. Springer Verlag, 2006.
- [17] Martin Pelikan, Kumara Sastry, and David E. Goldberg. Multiobjective estimation of distribution algorithms. In Martin Pelikan, Kumara Sastry, and Erick Cantú-Paz, editors, *Scalable Optimization via Probabilistic Modeling: From Algorithms to Applications*, Studies in Computational Intelligence, pages 223–248. Springer Verlag, 2006.
- [18] Peter AN Bosman and Dirk Thierens. The naive midea: A baseline multi-objective ea. In *Evolutionary Multi-Criterion Optimization*, pages 428–442. Springer, 2005.
- [19] N. Hansen, S.D. Muller, and P. Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, 11(1):1–18, 2003.
- [20] Hans-Georg Beyer and Hans-Paul Schwefel. Evolution strategies — A comprehensive introduction. *Natural Computing: an international journal*, 1(1):3–52, 2002.
- [21] Christian Igel, Nikolaus Hansen, and Stefan Roth. Covariance matrix adaptation for multi-objective optimization. *Evolutionary computation*, 15(1):1–28, 2007.
- [22] Chang Wook Ahn and R. S. Ramakrishna. Multiobjective real-coded Bayesian optimization algorithm revisited: Diversity preservation. In *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 593–600, New York, NY, USA, 2007. ACM Press.
- [23] Bo Yuan and Marcus Gallagher. On the importance of diversity maintenance in estimation of distribution algorithms. In *GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 719–726, New York, NY, USA, 2005. ACM Press.
- [24] Jonathan Shapiro. Diversity loss in general estimation of distribution algorithms. In *Parallel Problem Solving from Nature - PPSN IX*, pages 92–101, 2006.
- [25] Luis Martí. *Scalable Multi-Objective Optimization*. PhD thesis, Departamento de Informática, Universidad Carlos III de Madrid, Colmenarejo, Spain, 2011.
- [26] Peter A.N. Bosman and Dirk Thierens. Adaptive variance scaling in continuous multi-objective estimation-of-distribution algorithms. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation - GECCO '07*, page 500, New York, New York, USA, 2007. ACM Press.
- [27] Peter A. N. Bosman. The anticipated mean shift and cluster registration in mixture-based EDAs for multi-objective optimization. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation - GECCO '10*, page 351, New York, New York, USA, 2010. ACM Press.
- [28] Luis Martí, Jesús García, Antonio Berlanga, and José Manuel Molina. Introducing moneda: Scalable multiobjective optimization with a neural estimation of distribution algorithm. In *Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation*, pages 689–696. ACM, 2008.
- [29] Luis Martí, Jesús García, Antonio Berlanga, Carlos A. Coello Coello, and José Manuel Molina. MB-GNG: Addressing drawbacks in multi-objective optimization estimation of distribution algorithms. *Operations Research Letters*, 39(2):150–154, 2011.
- [30] Stephen Grossberg. *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control*. Reidel, Boston, 1982.
- [31] Luis Martí, Jesús García, Antonio Berlanga, and José M Molina. Multi-objective optimization with an adaptive resonance theory-based estimation of distribution algorithm: a comparative study. In *Learning and Intelligent Optimization*, pages 458–472. Springer, 2011.

- [32] Jüergen Branke, Clemens Lode, and Jonathan L. Shapiro. Addressing sampling errors and diversity loss in umda. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation — GECCO '07*, pages 508–515. ACM Press, 2007.
- [33] H. Mühlenbein and G. Paaß. From recombination of genes to the estimation of distributions I. Binary parameters. In Hans-Michael Voigt, Werner Ebeling, Ingo Rechenberg, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature - PPSN IV*, pages 178–187, Berlin, 1996. Springer Verlag. LNCS 1141.
- [34] J. M. Pena, Victor Robles, Pedro Larrañaga, V. Herves, F. Rosales, and M. S. Pérez. GA-EDA: Hybrid evolutionary algorithm using genetic and estimation of distribution algorithms. In *Innovations in Applied Artificial Intelligence*, pages 361–371, Heidelberg/Berlin, 2004. Springer.
- [35] Q. Zhang, Jianyong Sun, and Edward Tsang. An evolutionary algorithm with guided mutation for the maximum clique problem. *IEEE Transactions on Evolutionary Computation*, 9(2):192–200, April 2005.
- [36] Juan J Durillo and Antonio J Nebro. jmetal: A java framework for multi-objective optimization. *Advances in Engineering Software*, 42(10):760–771, 2011.
- [37] Simon Huband, Philip Hingston, Luigi Barone, and Lyndon While. A review of multiobjective test problems and a scalable test problem toolkit. *Evolutionary Computation, IEEE Transactions on*, 10(5):477–506, 2006.
- [38] Nicola Beume, Boris Naujoks, and Michael Emmerich. Sms-emoa: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3):1653–1669, 2007.
- [39] Saku Kukkonen and Kalyanmoy Deb. Improved pruning of non-dominated solutions based on crowding distance for bi-objective optimization problems. In *Proceedings of the World Congress on Computational Intelligence (WCCI-2006)(IEEE Press)*. Vancouver, Canada, pages 1179–1186, 2006.
- [40] G. Rudolph, H. Trautmann, and O. Schütze. Homogene Approximation der Paretofront bei mehrkriteriellen Kontrollproblemen. *Automatisierungstechnik (at)*, 60(10):612–621, 2012.